CDF/DOC/COMP_UPG/CDFR/XXXX
version 1.3 December 15, 2005

Proposal for a dCache-based Analysis Disk Pool for CDF (DRAFT)

CDF Collaboration

Abstract

This note presents the motivation and proposed deployment plan for an analysis disk pool based on dCache.

Contents

1 Motivation

At present, CDF maintains approximately 120 TB of disk that is managed by the physics groups for the purpose of supporting on-going analyses. The two most important roles played by this so-called ”project disk” include the use as a staging area for the collection, concatenation, validation and uploading to tape of production secondary and ntuple datasets that are produced and maintained by the physics groups; and as a storage area for analysis-specific data and ntuples, all of which have a useful lifetime limited to the publication cycle for the particular analyses. Data in the latter category is not stored on tape. Historically, the production ntuple datasets were also stored and served from the project disk. This use case is declining as these datasets become managed by the Data Handling (DH) system. Data in the project disk are served to desktops or the CAF primarily via rootd.

The current system has several deficiencies. First, effective management and efficient use of the project space is hampered by the large number of service level administrators and the distribution of allocated space across many fileservers. Second, data stored in the pool are cataloged by individuals and on web pages maintained by the physics groups. Re-allocation or maintenance of resources within the pool invariably causes disruption to the users and catalog maintainers. Finally, there is no mechanism to balance the load across the system.

2 Proposed solution

CDF proposes to replace a part of the static project disk by a collection of fileservers served by dCache. The use of dCache offers the following features:

- A product support by Fermilab, with considerable experience and knowledge base at the lab and at CDF. Other experiments within the HEP community, including the LHC experiments, have adopted dCache as their file caching and serving technology.

- A global namespace that will simplify organization and cataloging of data, as well as monitoring and accounting of resource utilization by system administrators.

- Virtualization of the disk space. By de-coupling the namespace from the physical devices, maintenance of individual fileservers and re-allocation of resources can in many cases occur without impact to users or catalog maintainers.

- A highly scalable and configurable file serving capability.
- An existing tool set that provides system-wide monitoring capabilities.
- A system for which the analysis software of the experiment is already designed to use.

The established features and capabilities of dCache make it possible to deploy the system with a minimum of development effort by the experiment. With dCache, CDF can transform a substantial fraction of the project disk from a collection of independently organized directories to a centrally managed resource.

The system will support the following use cases:

- Provide space for the collection, concatenation and validation of production datasets and ntuples maintained and produced by the physics groups. After validation, these ntuples will be stored in the DH system.
- Provide buffer space for ntuples and MC production prior to uploading to the DH system.
- Provide temporary storage for secondary and tertiary datasets or ntuples produced by the physics groups for specific analyses or studies.
- Temporarily host smaller analysis-specific datasets and ntuples produced by individual or groups of users.
- Serve all of the above data to desktops and the CAF.

The data associated with specific analyses are not considered as candidates for archiving in the DH system, and will be deleted after publication of the associated analyses. None of the data within the disk pool will be backed up to tape within the DH system; it is the responsibility of the users to back up files as needed by some other means. Current policy requires a justification for files to be stored in the central tape facility, and recommends use of RAID-based disk to minimize risk to disk-resident files.

Note that production ntuples will reside and be served from the disk pool, just as they currently reside on static disk, during the transition phase to an end-to-end SAM solution for ntuple analysis. This use case may well disappear by the time the pool goes into production.

3 Risks

By moving from an intrinsically distributed system to a central service, all activity within the pool becomes susceptible to weaknesses or instabilities of the central service. Limitations in the performance of pnfs at high transaction rates, for instance, place limits on the overall performance of the system. Misuse or accidents by a small number of users can degrade the performance across the entire pool.

The use of dCache to serve ntuples to the CAF and to interactive desktop users represents a new use case for dCache, one for which the experiment has limited experience. Ntuple analysis programs typically read files at a much higher rate than do those that read production output files, which increases the relative transaction rate required of dCache.

Finally, user expectations of the system may not match the capabilities of the system or the supported use cases. The namespace presented by dCache, for instance, is structured like that of a filesystem. Some commands, such as ”ls”, are provided as well. Such features can result in misconceptions about the capabilities of the system, which can lead to misuse.

The priority of the prototyping and deployment plan will be to establish stability of the overall system and to minimize the operational workload. To this end, we have adopted the following strategy to mitigate the above risks. First, the prototyping phase will carefully study and characterize the supported use cases. The existing prototyping effort has already accumulated considerable experience with production ntuple collection and validation – these are easily supported by the existing prototype system. Concatenation and uploading are centralized activities, and therefore do not represent significant additional loads on the system. The remaining use cases will be studied by monitoring the system under operating conditions with a small number of test volunteers.

Second, the overall performance capabilities of the system will be mapped using tests designed to stress specific features or parameters. Monitoring during these tests will allow identification of the limiting elements under various conditions and extremes.

Third, careful monitoring of the system as the number of users and system capacity is scaled up will ensure that performance metrics remain within the stable operating region established by system testing. It is also important during this phase to confirm that the load follows expectations based upon the test user experience.

Fourth, the deployment team will establish a set of user guidelines that will explain the supported use cases, best practices and proscribed activities. Where possible, technical means will be used prevent or reduce the possibility of misuse, such as limiting the number of nodes with direct pnfs mount points.

4 Support model

CDF is committed to providing the long-term operational and service level support for the analysis disk pool, and intends to request only a minimum of effort from the Computing Division. While many details of the support to be requested from the CD have not yet been determined, pending further discussion with the relevant CD groups, we can describe the general features of the proposed support model.

In broad terms, we propose a three-tiered support model. The first level of support will be provided by “power users” and the CDF operations group. This group will perform daily monitoring of the system and will triage service level problems as they occur. They will use tools that allow them to identify the cause of the most common problems and possess a limited capability to re-start certain non-privileged components should the need arise.

The second level of support will be the responsibility of a CDF institution and will be formalized in an MOU agreement between the experiment and the institution. Individuals involved at this level will require sufficient expertise to be able to identify less common problems and solutions, and to be able to recognize issues indicative of possible configuration problems. The second tier maintainers must also take responsibility for implementing any required changes in the configuration of the system, as well as handling routine tasks such as adding or removing fileservers and re-aligning pools. Second tier support personnel will be the point of contact between the experiement and the CD on issues regarding the disk pool. The University of Michigan currently fills the second tier support role.

The proposed third tier of service-level support will come from the CD. Initially, CDF requests only expert-level consulting within negotiated limits. It is hoped that this will be sufficient to carry the project from prototyping through stable production running. No development support is requested since the pool relies upon the existing capabilities and feature set.

Hardware and OS support will continue to come from the REX system administration group, as it does now for the existing project disk.

Finally, we request the possibility of re-examining the service-level support agreement between the experiment and the CD after six months to one year of stable production operation. This period will allow the experiment and the CD to understand in detail the level of effort required to operate the system. Should the parties agree that level of second-tier effort is sufficiently small, then CDF may request that certain of those activities be consolidated in the existing dCache/ISA or REX dCache support personnel.

5 Outline of the deployment plan

The following outline describes the specific steps, phases and overall time scales of the three-stage deployment plan. More details of the plan, including a resource-loaded schedule, are provided in appendices.

6 Phase I: prototyping, testing and understanding performance

Characteristics of the system (started in summer 2005).

Size of the system: up to 50 TB (= 25 TB for ntuple production + 5 TB * number of physics groups participating)

6.1 A. Project goals

1. Develop a resource loaded schedule for the deployment with deliverables and milestones. This schedule will be maintained to track progress through the prototyping and preproduction. (See Appendix A for this schedule.)

2. Gain understanding of the technical characteristics and performance parameters of the system.

3. Gain a detailed understanding of the supported CDF use cases and system requirements.

4. Discuss with CMS and CCF the possibility of building a mutual knowledge base through meetings to talk about issues of common interest.

5. Use the experience and feedback of the power users to refine the rules and usage guidelines (such as identifying good and bad use patterns). This version of the guidelines will be suitable for distribution to general CDF users in later deployment phases, and will form the basis of recommendations to the collaboration.

6. Develop specifications for the production hardware. (See Appendix B.)

7. Develop support agreements for Phase II.

   - Support during this Phase II will come primarily from the CDF collaboration.
   - Request expert-level consultant support from CD/CCF.

6.2 System characterization steps

1. Determine performance characteristics of the Disk Pool system by performing various tests of functionality, load and scalability. (See Appendix C for details of the tests and metrics.)

   a) Document the results of the tests and communicate results with dCache developers.

   b) Determine the range of operating parameters that define stable operation of the system.

2. Investigate and understand failure modes of the system.

3. Develop an automatic monitoring system for the Disk Pool. The monitoring used for the Disk Pool will be based upon that already developed for the production dCache system. In this way CDF will use common, centrally supported tools and will update both dCache systems' monitoring as improvements become available.

6.3 Characterization of supported use cases

1. Recruit experienced power users willing to use the Disk Pool for specified use cases. The understanding with these users will acknowledge that the system is not yet in production. Interruptions due to system instabilities or tests may occur on a regular basis. These users must also agree to work with the CDF offline group to help with system characterization.

2. Develop an initial set of Disk Pool usage guidelines for the power users. Among other items, these guidelines will:

   - define minimum file size requirements (1GB)

- prohibit the use of tar, unzip, etc., within the Disk Pool
- prohibit compilation, etc., within the Disk Pool
- prohibit the use of "ls" on the pnfs namespace
- prohibit the storage of log files to the Disk Pool, and provide procedures to direct log files to the static disk of the CDF interactive login pool

(See Appendix E for the proposed guidelines.)

3. Train power users to monitor the Disk Pool performance, to use the monitoring to diagnose its failure modes and to perform limited administration tasks (such as re-starting doors or placing Helpdesk calls).

4. Monitor the system under load from the power users to estimate how the system load will scale with additional users.

6.4 D. Pre-production readiness review

1) At the end of Phase I, hold an internal preproduction readiness review to review the progress ahead of the transition between prototyping and preproduction. Members of the FNAL Computing Divison should be invited to attend the review in order maintain good communication.

6.5 E. Timescales and accomplishments for Phase I

- CDF offline group plans to complete Phase I by the end of January 2006. By that time:

  a) Performance of the system under a specified user load is understood. Total performance requirements for the system are specified:

  - served data volume: XXX TB/day
  - aggregate bandwidth provided by the system: XXX GB/sec
  - max supported transaction rate: XXX/sec
  - max supported number of the concurrent users: XXX

- The capabilities of the system are understood.

- Basic monitoring system is in place.
- Experience of the power users is documented.
- Production hardware specification completed.
- Ready for pre-production review.

7 Phase II: pre-production

Size of the system: transitions from prototype size to full production size (up to 100 TB)

7.1 A. Project goals

1. Verify user load models developed during prototyping phase.

2. Transition to production operational mode.

3. Demonstrate sustained stable operations at full production load.

4. Develop long-term support agreements for Phase III.

   - Develop agreements with Univ. of Michigan and CDF collaborators to establish a two-tiered operational model
     - first level provided by power users and CDF operations group
     - second level provided by Univ. of Michigan group
   - Request expert-level consultant support from CD/CCF

7.2 B. Transition to production operational mode

1. Expand user base of the system incrementally and verify user load models. User agreement will specify that the system is not yet in production. Interruptions due to system instabilities or tests may occur on a regular basis.

   a) Use authentication mechanism to limit access.

   b) Monitor system performance and load, and compare with expectations.

   c) Revise usage policies or configuration as needed to ensure load remains within operational limits.

2. Transition to production hardware and full production capacity.

   a) Repeat load tests after each transition or expansion to determine operational limits.

3. Institute production operational policies.

   a) Establish fully automated monitoring

b) Institute policies to minimize operational load

c) Establish stable operation with sparse, regularly scheduled maintenance downtimes.

4. Expand user base to full collaboration under pre-production user agreement.

7.3 C. Production readiness review

1. Convene and complete a production readiness review jointly with the Computing Division. The charge for this review will be determined jointly between CDF and the CD. Satisfactory completion of this review is required before the system can be declared in production.

7.4 D. Schedule and accomplishments for Phase II

Deployment Phase II is expected to take between 2 and 4 months. At the end of this time:

- The system will be operating stably at full load with production hardware.
- Long-term support agreements are in place.
- Ready for production review.

8 Phase III: production operation

Size of system: up to about 100 TB

Project goals

1. Maintain stable operations of the system using the production operational policies and support agreements established during Phase II.
2. Add or re-allocate resources as needed to meet the needs of the supported use cases.
    - Perform tests as needed during the regularly scheduled downtimes to validate configuration and policy changes.

A Resource-loaded schedule

B Production hardware specification

C System tests and metrics

In this appendix, we describe the tests we plan to run on the diskpool prototype. We envision two kinds of tests:

- tests mimicking user activity that are designed to report the performance that users will expect to see under different load conditions;

- system monitoring tests that report on the relevant metrics of pnfs and dCache.

C.1 Tests mimicking user activity

These tests are intended to map the performance of the system over a range of loading conditions. An important function of the tests will be to measure how the performance of the system varies over the expected range of user loads, possibly identifying some conditions under which the system breaks.

C.1.1 Test 1

This test is designed to report the data transfer rates to be expected when some users are reading data and others are concurrently writing data into the diskpool. Approximately 100 users run on the CAF on a typical day, with 400 registered users on the CAF. Most of the users are analysers reading data. A typical user spends most of his/her analysis time reading data, and only a small fraction generating data. Based on these numbers, we propose the following test:

Concurrent transactions: 100, 10, 5, 0 READS coupled with 50, 10, 5, 0 WRITES starting at the same time. Participating files are all unique. Present the average and total measured data rates (in units of MB/sec), as well as the data rates for the fastest and slowest clients. For example, the total for 50 writes is (50 * 4GB) / (max time − start time). Here 'max time' is the time taken by the slowest client. For this example, the average is equal to (total / 50). Report these results for the WRITES and the READS.

C.1.2 Test 2

The CAF has ≈ 4500 VMs. We know from past monitoring that up to about 30% of the CAF is used for ntuple analysis. Some fraction of this 30% will be used to analyse data off the diskpool. Based on this observation, we propose the following test:

Have N clients, where N can be as large as 1500, access the ROOT files in diskpool. All clients start at the same time. Each client does TFile::Open("dcap://.....") on 30 unique files one by one. About $N \times 30$ files already produced by users will be used for this test. Record the smallest, largest and average file size in this test. Record the file transaction frequency, defined as (30 files) / (delta Time). Show a histogram of the frequencies reported by the N clients. Also report the average file transaction frequency for the entire N-client job, defined by $(N * 30)$ / (first client start time $-$ last client end time).

D System Monitoring

D.1 Operational experience

- On Nov 30 2005 we have identified a problem which turned out to be the major source of Disk Pool door hang-ups - breakdowns in the interactions with the kerberos server. Since it has been fixed (round-robin of 6 kerberos servers, thanks to A.Kulyavtsev/V.Podstavkov form CD/CCF) the DiskPool was running smoothly.
- Dec 16 2005 - one of the fileservers (fcdfdata29) developed a hardware problem, fixed by the REX sysadmin group (?) within several hours . Exact source: interrogate Alexei.

Plot in Fig. 1 shows distribution for the number of client processes simultaneously accessing the Disk Pool as a function of time.

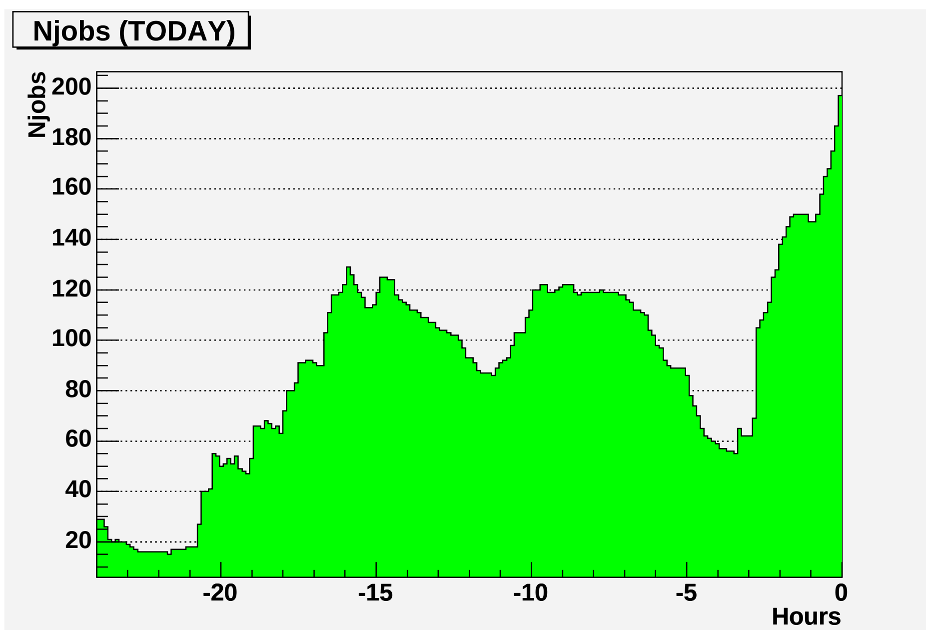

Figure 1: number of client processes simultaneously accessing the Disk Pool vs time

Plot in Fig. 2 shows distribution for the I/O rate per process as a function of time within the last 24 hours.

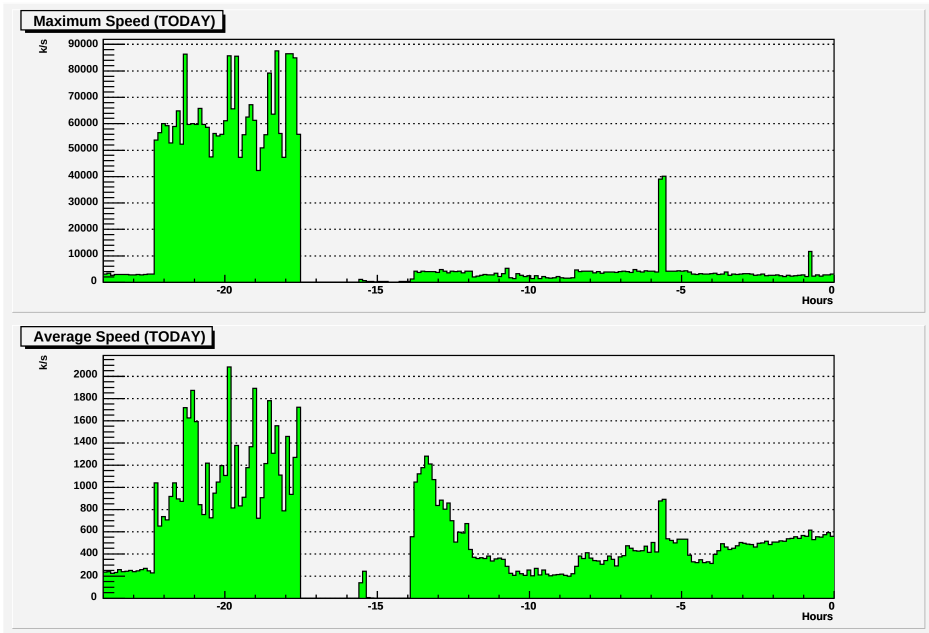

Figure 2: number of client processes simultaneously accessing the Disk Pool vs time

Plot in Fig. 3 shows distribution for the number of clients as a function of time durint the load test (2005-12-19)

E Proposed usage guidelines

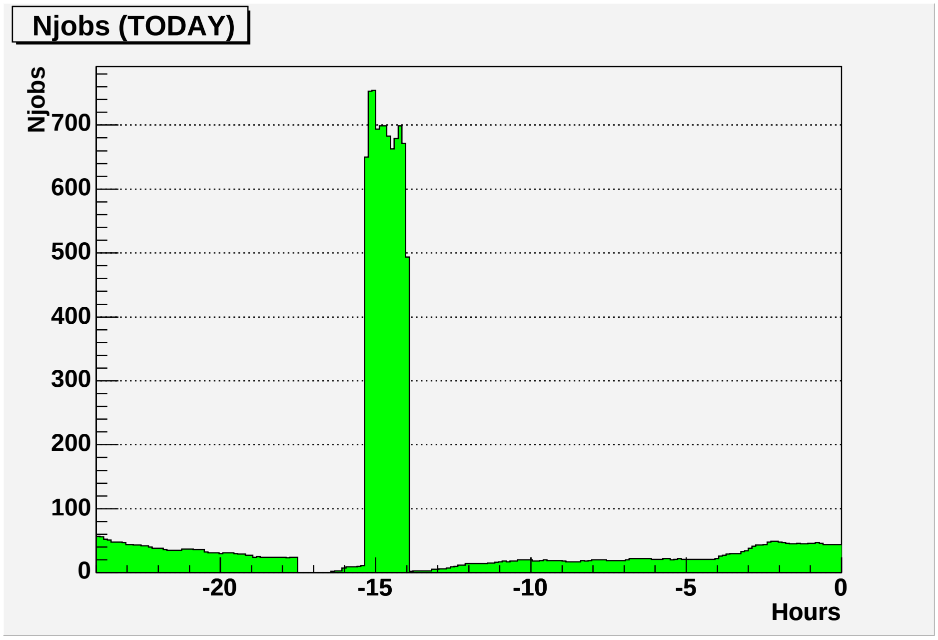

Figure 3: number of client processes simultaneously accessing the Disk Pool vs time during the load test 2005-12-19